Efficient Data-Driven Reference Governor Design for Safe Evasive Manoeuvring

Petar Velchev, Alberto Bertipaglia, Felipe Santafe, Mohammad Khosravi, and Barys Shyrokau

***Abstract*—This paper presents a novel data-driven Reference Governor with Model Predictive Control, integrating local motion replanning and path following for collision avoidance. Employing a model-free Reference Governor, the proposed framework utilises system knowledge through Bayesian Optimisation to augment predetermined evasive trajectories, minimising path-following errors and simultaneously ensuring obstacle safety margins. A single-track vehicle model in combination with a non-linear tyre model is used to capture the vehicle's dynamics. The optimised control action is the vehicle steering angle, whilst the Reference Governor optimises parameters of a sigmoid reference signal to minimise the tracking error and guarantee safety with respect to obstacles in emergency manoeuvres. The proposed approach is evaluated on a single lane change using a high-fidelity simulation environment, and its performance is compared to a baseline controller integrating path following and obstacle avoidance. The results demonstrate a 14% reduction in safety critical overshoot, maximising obstacle safety distance and a four times lower controller cycle time compared to the baseline. Furthermore, through a robustness analysis, it is demonstrated that the proposed approach is more robust towards model mismatches and perception-based errors, as seen by average 30% and 40% reductions in near-miss and collision rates.**

***Index Terms*—Obstacle avoidance, model predictive control, reference governor, path following, Bayesian optimisation.**

I. Introduction

AUTOMATED vehicles operating near the handling limits must be capable of executing evasive manoeuvres with high precision and robustness to avoid collisions in dynamic and uncertain environments. In such scenarios, due to high non-linearities, deviations from a reference path can result in safety-critical situations, making the design of accurate and adaptable path following systems essential [1], [2].

Received 31 December 2025; accepted 23 May 2026. Date of publication 15 June 2026; date of current version 8 July 2026. This article was recommended for publication by Associate Editor T. Berrueta and Editor S.-J. Chung upon evaluation of the reviewers' comments. This work was supported by MOCO Project, funded by European Union, under Grant agreement 101183051. *(Corresponding author: Petar Velchev.)*

Petar Velchev is with Capgemini Engineering, 1831 Machelen, Belgium, and also with the Delft Center for Systems and Control, Delft University of Technology, 2628 CD Delft, The Netherlands (e-mail: petar.velchev@capgemini.com).

Alberto Bertipaglia and Barys Shyrokau are with the Department of Cognitive Robotics, Delft University of Technology, 2628 CD Delft, The Netherlands (e-mail: A.Bertipaglia@tudelft.nl; b.shyrokau@tudelft.nl).

Felipe Santafe is with Toyota Motor, 1930 Zaventem, Belgium (e-mail: felipe.santafe@toyota-europe.com).

Mohammad Khosravi is with the Delft Center for Systems and Control, Delft University of Technology, 2628 CD Delft, The Netherlands (e-mail: mohammad.khosravi@tudelft.nl).

Digital Object Identifier 10.1109/LRA.2026.3703595

A key challenge is the ability to locally and rapidly re-plan trajectories while maintaining safe distances from obstacles. This is critical for advanced driver assistance and autonomous driving systems, particularly in emergencies requiring fast, reactive decisions near the vehicle's handling limits. Applications include urban scenarios with unpredictable pedestrians, highway cut-in scenarios, and evasive manoeuvres. However, re-planning and path following at the limit of handling is inherently challenging due to the highly non-linear vehicle dynamics and the need to balance tracking performance with safety constraints in real-time [3], [4]. Classical motion planning and trajectory tracking methods assume linear dynamics through the use of basic point mass models, which fail to capture the complex interactions between control inputs, tyre forces, and vehicle responses under high slip or load transfer conditions [5], [6]. Furthermore, incorporating real-time re-planning into a control pipeline with minimal computational burden remains a persistent difficulty [7], requiring dedicated hardware instead of implementing on existing less powerful in-vehicle control units.

To address these challenges, this paper proposes a novel control architecture combining a model predictive controller (MPC) with a model-free Bayesian Optimisation-based Reference Governor (BORG) [8]. The MPC is responsible for tracking a reference trajectory by computing the optimal steering input, using a non-linear single-track vehicle model with non-linear tyre dynamics to balance model complexity and computational efficiency. The reference governor (RG) optimises the parameters of a sigmoid-shaped evasive reference trajectory, visualised in Fig. 1 using BO. This decoupled architecture enables the system to adaptively adjust the planned manoeuvre in real-time, reducing tracking error by compensating for model mismatch and improving vehicle to obstacle distance while maintaining computational efficiency.

II. Related Work

Efficient obstacle avoidance in automated vehicles often relies on simplified models such as kinematic single-track or point-mass approximations, favoured for their low computational cost and real-time suitability in structured or low-speed scenarios [5], [6]. However, while effective in the linear regime, these models fail to capture key dynamics such as tyre slip, load transfer, and actuator limits, which are critical in aggressive or high-speed manoeuvres. This oversimplification [9] can lead to unsafe or infeasible trajectories due to poor planning-control consistency, while neglecting high-fidelity dynamics [10] often results in

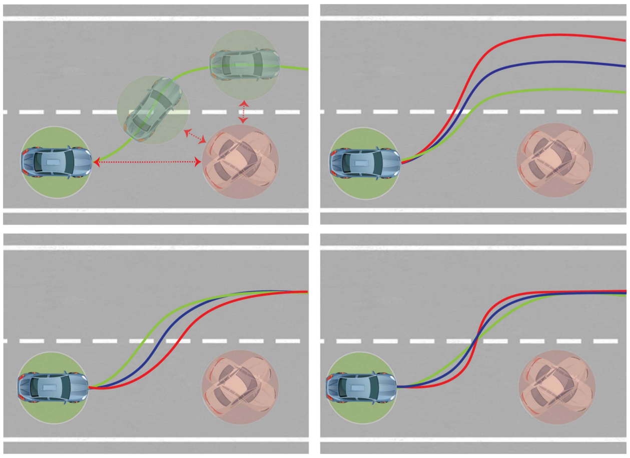

Fig. 1. Top left diagram depicts vehicle to obstacle distance (V2O) for a target manoeuvre in green. The top right, bottom left and bottom right show a variation in end location, starting point and steepness of manoeuvre.

overly conservative or suboptimal paths. Recent work [11] has therefore shifted toward hybrid approaches that combine accurate dynamic models or learning-based surrogates with simplified planners to improve robustness while maintaining real-time feasibility. In parallel, data-driven methods have been applied for predictive model correction [12], [13] and controller tuning [14], [15], [16], [17]. Reference Governors (RGs) modify reference inputs to enforce constraints without full re-optimisation. Originally proposed for linear systems [18], they have been extended to vehicle steering [19], Lyapunov-based constraint enforcement [20], and manoeuvre-level decision-making via reachability analysis [21]. RGs can also act as MPC pre-filters, reducing constraint violations and computational load [22], while recent learning-based variants derive safe sets from data to improve adaptability in uncertain environments [23]. In summary, simplified re-planners lack fidelity near handling limits, while traditional RGs often rely on accurate models or lack adaptability. Hybrid and data-driven approaches [24], [25] offer promising directions for robust, real-time control in automated driving.

Motivated by the limitations of existing approaches, the proposed BORG architecture introduces a novel decoupled framework for evasive manoeuvring that separates trajectory adaptation from control. Unlike conventional MPC formulations that embed obstacle avoidance directly within the cost function leading to increased computational burden and sensitivity to model inaccuracies, the proposed method employs a model-free Bayesian Optimisation-based Reference Governor (RG) to adapt the reference trajectory prior to control execution.

The key novelty lies in (i) formulating obstacle avoidance as a reference augmentation problem using a low-dimensional parametric trajectory representation, and (ii) solving this problem offline via Bayesian Optimisation to learn a surrogate Gaussian Process (GP) model that enables fast online adaptation. This results in a single-shot reference update, avoiding iterative re-optimisation during runtime while preserving responsiveness to changing conditions.

The RG, depicted in Fig. 2, takes as input the nominal lane change reference, vehicle states, and obstacle information, and

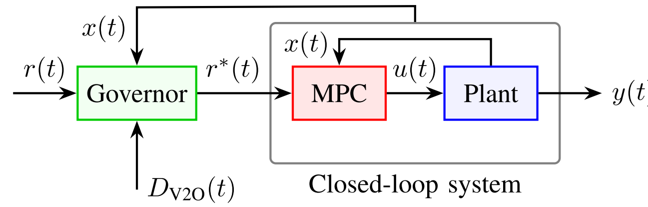

Fig. 2. Block diagram of the reference governor [19] applied to a closed-loop MPC system for obstacle avoidance.

outputs an augmented trajectory defined by optimised sigmoid parameters. This augmented reference is then tracked by a standard MPC, which remains focused on control execution rather than re-planning.

The contributions of this paper are threefold:

1) A novel decoupled RG–MPC architecture, where a data-driven Bayesian optimisation-based RG performs real-time trajectory adaptation via a pre-trained GP surrogate, reducing online computational complexity compared to integrated MPC formulations.
2) A parametric reference design that explicitly accounts for handling limits, enabling smoother and less aggressive manoeuvres, achieving a $14\,\%$ reduction in overshoot relative to a baseline.
3) Improved robustness to model mismatch and perception uncertainty through model-free reference adaptation, demonstrated by average reductions of $30\,\%$ in collision rate and $40\,\%$ in near-miss rate across varying velocities.

III. Reference Governor Design Problem

During an evasive manoeuvre, the path following algorithm is responsible to track the reference trajectory, satisfying the system and operational constraints, e.g., the actuator physical limits. Nonetheless, if the reference is excessively aggressive, it may lead to infeasible or unsafe behaviour. Accordingly, we need to develop an RG, as shown in Fig. 2, to modify the reference in real-time based on the vehicle state and constraints, i.e., the RG is required to provide an optimal feasible reference r^*, ensuring safe and constraint-compliant tracking by the controller. Consider a parametric class of reference target signals $\{r_\theta \mid \theta \in \Theta\}$, each parametrised by a vector $\theta \in \Theta$, where $\Theta \subseteq \mathbb{R}^{n_\theta}$ denotes the set of admissible parameters. Let θ^* be the optimal parameter vector that characterises the desired target reference trajectory r^* as r_{θ^*}. More precisely, to obtain r^*, the parameters should be configured to optimise the overall closed-loop performance while ensuring the safety of the system and desired tracking behaviour. By analysing the tracking error, which represents the difference between the reference and optimised path, various metrics can be introduced to assess tracking quality, such as overshoot magnitude, settling time, steady-state error, rise time and other application specific indicators. Given a fixed control structure, the tracking quality and errors depend mainly on the vector of parameters θ. Accordingly, each of the above performance metrics is essentially a function of θ, denoted by $c_i : \Theta \to \mathbb{R}$, for $i = 1, \ldots, n_c$. Thus, the *overall performance metric*, $f : \Theta \to \mathbb{R}$, can be defined as a weighted sum of the individual performance metrics, i.e., we have

$$f(\theta) := w^{\mathsf{T}}c(\theta) = \sum_{i=1}^{n_c} w_i c_i(\theta), \quad \theta \in \Theta, \tag{1}$$

where $c(\theta) := [c_i(\theta)]_{i=1}^{n_c}$ and $w := [w_i]_{i=1}^{n_c}$. The weights $w_1, \ldots, w_{n_c}$ are determined based on the relevance and scale of the associated performance metrics. Given the defined performance metrics, we can formalise the reference augmentation problem. Specifically, to find the optimal reference parameters θ^* while maintaining safety, we need to solve the optimisation problem

$$\min_{\theta \in \Theta} f(\theta) = w^{\mathsf{T}}c(\theta). \tag{2}$$

The objective function f is not analytically computable, i.e., it lacks tractable closed-form expressions, even if the system dynamics are well understood and accurately known. On the other hand, for each reference parameter $\theta \in \Theta$, an experiment can be performed on the system to measure the tracking error signal and subsequently evaluate $f(\theta)$. Therefore, these functions are accessible only through a *black-box oracle*. To augment the reference parameters according to the optimisation problem above, one needs to employ a suitable data-driven optimisation method where the objective function is represented as expensive-to-evaluate black-box functions. To this end, we propose a reference governor design based on Bayesian optimisation (BO), as described in the following sections, and demonstrate its application through integration with an MPC vehicle motion control scheme.

IV. Governed Vehicle Motion Control System Design

The proposed reference governor for a model predictive controller as depicted in Fig. 2 to meet tracking and obstacle safety performance goals. This section first describes the predictive model, the setup of the optimal control problem, the Bayesian optimisation formulation and lastly the utilisation of the RG. The latter is split into two phases:

- Offline: RG is trained using BO through the use of GP surrogate modelling to augment reference parameters.
- Online: the trained GP surrogate model is used as a reference augmentation filter in an online fashion.

A. Vehicle Predictive Model

The vehicle dynamics are modelled using the non-linear single track model, which captures the lateral, yaw, and longitudinal behaviour of the vehicle in a planar environment. The model assumes a single front and rear wheel to simplify the analysis while preserving key dynamic characteristics. The system states include the longitudinal and lateral velocities respectively denoted by v_x and v_y, yaw rate r, yaw angle ψ, global position coordinates (x_p, y_p), and steering angle δ. Tyre lateral forces are represented using a linear tyre model with cornering stiffness coefficients C_{α_f} and C_{α_r} [26]. The resulting equations of motion describe the time evolution of the vehicle's dynamics under the influence of steering inputs and inertial properties as

$$\begin{aligned}
&\dot{v}_x = v_y r, \quad \dot{\psi} = r, \quad \dot{\delta} = d_\delta, \\
&\dot{v}_y = -\frac{C_{\alpha_f} + C_{\alpha_r}}{m v_x} v_y + \frac{L_r C_{\alpha_r} r - L_f C_{\alpha_f}}{m v_x} - v_x r + \frac{C_{\alpha_f}}{m}\delta, \\
&\dot{r} = \frac{L_r C_{\alpha_r} - L_f C_{\alpha_f}}{I_z v_x} v_y - \frac{L_r^2 C_{\alpha_r} + L_f^2 C_{\alpha_f}}{I_z v_x} r + \frac{L_f C_{\alpha_f}}{I_z}\delta, \\
&\dot{x}_p = v_x \cos(\psi) - v_y \sin(\psi), \quad \dot{y}_p = v_x \sin(\psi) + v_y \cos(\psi).
\end{aligned} \tag{3}$$

The cornering stiffnesses C_{α_f} and C_{α_r} are approximations of the tyre lateral force response at small slip angles, derived from the *Magic Formula* tyre model [27], that describes the non-linear relationship between tyre slip angle and lateral force, typically expressed as $F_y = D\sin(C\arctan(B\alpha - E(B\alpha - \arctan(B\alpha))))$, where α is the slip angle, and B, C, D, and E are empirical fitting factors. The cornering stiffness is obtained through this relationship, yielding an approximation of the tyre behaviour in the small-slip regime, allowing the non-linear single track model-based controller to remain computationally efficient while retaining physical accuracy. Additionally, the Dugoff tyre model is also utilised in the predictive model in combination with a high-fidelity vehicle model [28]. This system is used to evaluate the performance of the proposed BORG with a still computationally efficient MPC for real-time implementation but in combination with a very accurate vehicle simulation as a plant.

B. Optimal Control Problem

The presented optimal control problem is formulated for use within a MPC framework to enable safe and smooth trajectory tracking with obstacle avoidance. Accordingly, given suitable sampling time t_s and prediction horizon N, the baseline MPC cost function J is defined as

$$J = \sum_{i=1}^{N} \Bigg[q_{e_{Y_p}} e_{Yp,i}^2 + q_{e_{X_p}} e_{Xp,i}^2 + q_{\dot{\delta}} \dot{\delta}_i^2 + q_{\dot{F}_x} \dot{F}_{x_i}^2 + {} + \sum_{j=1}^{N_{\text{obs}}} \left(q_{e_{\text{V2O}}} e_{\text{V2O},j,i}^2 \right) \Bigg], \tag{4}$$

where $e_{Yp,i}$ and $e_{Xp,i}$ represent the lateral and longitudinal position tracking errors at time step i, weighted by $q_{e_{Y_p}}$ and $q_{e_{X_p}}$, respectively, $\dot{\delta}_i$ is the steering rate, penalized through $q_{\dot{\delta}}$ to ensure smooth steering actions [29], $\dot{F}_{x_i}$ denotes the longitudinal force rate, penalized by $q_{\dot{F}_x}$ to prevent aggressive acceleration or braking, and, $e_{\text{V2O},j,i}$ is the distance-based error term representing proximity to the j-th obstacle at time step i, weighted by $q_{e_{\text{V2O}}}$, to enforce obstacle avoidance.

The obstacle distance metric is defined as

$$D_{\text{V2O}} = \left((X - X_{\text{obs}})^2 + (Y - Y_{\text{obs}})^2 \right)^{\frac{1}{2}} - r_{\text{obs}} - r_{\text{veh}}, \tag{5}$$

which represents the Euclidean distance between the vehicle and the obstacle centres, reduced by their respective safety radii r_{obs} and r_{veh}. The inclusion of D_{V2O} encourages the baseline controller to maintain a safe buffer from detected obstacles [12]. The BORG system operates under the same objective function as the baseline (4), but without obstacle avoidance consideration within the MPC objective function.

To ensure vehicle stability, constraints are imposed on the side slip angle and its rate, bounded respectively within $\pm 5^\circ$ and $\pm 25^\circ$ /s, steering wheel angle and velocity, limited respectively to ± 2.76 turns and $\pm 800^\circ$ /s, and lateral acceleration bounded within $\pm 0.85\, \mu g$. The introduced constraint reflects the fact that the required yaw rate to achieve a lateral manoeuvre cannot always be achieved due to the road friction coefficient not being able to provide the required tyre forces. Therefore, the yaw rate is bounded by a function based on the friction coefficient through a safety factor [4]. Furthermore, the limits reflect physical actuator capabilities and define a stable operating envelope during aggressive manoeuvres [26].

C. Sigmoidal Reference Trajectory Parametrisation and Bayesian Optimisation

Let the base reference trajectory be defined as a lateral displacement through a function of the longitudinal position x. In the proposed design, the reference governor augments this lateral reference by a sigmoidal parametrisation defined as

$$r_\theta(x) = \vartheta_1 \left(1 + e^{-\vartheta_2\,(x-\vartheta_3)}\right)^{-1}, \tag{6}$$

where $\theta := [\vartheta_1, \vartheta_2, \vartheta_3]^\top$ is the vector of parameters characterising the reference trajectory. Note that ϑ_1, ϑ_2, and ϑ_3 respectively determine the final lateral offset, the spatial centre of the manoeuvre, and the steepness of the transition, indicating the aggressiveness of a lane change manoeuvre. These features are implicitly demonstrated in Fig. 5. Accordingly, the choice of parameter vector θ directly affects the ability of the controller to achieve accurate and smooth path following while maintaining obstacle avoidance. This class captures typical smooth, monotonic collision avoidance manoeuvres, while more complex trajectories could be represented by higher-dimensional parametrisations at the expense of increased optimisation complexity. The parametrisation law (6) defines a class of lateral reference trajectories $y_{\text{ref}}(x) = r_\theta(x)$ to be tracked by the MPC. The objective is to determine an optimal trajectory, suitable as a controller-friendly reference signal, by selecting the parameter vector θ^* that optimises closed-loop behaviour, thereby improving both path-following accuracy and obstacle-avoidance performance. To this end, we introduce a suitable metric that aggregates various performance indices, including tracking accuracy, control effort, and obstacle safety. The resulting cost function $f : \Theta \to \mathbb{R}$ is defined as

$$f(\theta) = w_J J_{\text{T}}(\theta) + w_{Yp}\bar{e}_{Yp}^2 + w_{Xp}\bar{e}_{Xp}^2, \tag{7}$$

where $\bar{e}_{Yp}$ and $\bar{e}_{Xp}$ are the root mean square of the lateral and longitudinal tracking errors of the driven BORG path with respect to the reference trajectory, and $J_{\text{T}}(\theta)$ is the sum of the MPC cost as defined in (4) for the total manoeuvre. The structure of the introduced objective function ensures that the trajectory, which corresponds to its minimum, produces a reference that minimises tracking error, maximises ego vehicle to obstacle distance, and additionally enforces that the augmented reference deviates minimally from the nominal reference. Accordingly, we need to solve the optimisation problem formulated as

$$\theta^* = \text{argmin}_{\theta\in\Theta} f(\theta), \tag{8}$$

to obtain optimal trajectory characterised by θ^*.

Following our discussion in Section III, we know that f is an expensive-to-evaluate black-box oracle. Therefore, to efficiently tune the parameters and obtain θ^*, it is required to employ a suitable Bayesian optimisation scheme. Thus, to solve (8), a GP surrogate model is iteratively trained to approximate f from sampled evaluations [30], and subsequently, the obtained surrogate model is utilised to select new candidate parameters by optimising an acquisition function $\alpha : \Theta \to \mathbb{R}$ that balances exploration of uncertain regions and exploitation of promising areas in the parameter space Θ. Specifically, one can use the *Expected Improvement* (EI) acquisition function defined as

$$\alpha_{\text{EI},m}(\theta) = \mathbb{E}[\max(0, f_{\min,m} - f_m(\theta))], \tag{9}$$

where m indicates the iteration index, $f_{\min,m}$ denotes the best observed value of the objective function up to the current iteration, and f_m is the learned surrogate function obtained through *Gaussian Process Regression* (GPR) [30]. More precisely, the corresponding surrogate GP model provides a posterior mean function $\mu_m : \Theta \to \mathbb{R}$ and posterior variance function $\sigma_m^2 : \Theta \to \mathbb{R}$, allowing the EI acquisition function to be computed in closed-form as

$$\begin{aligned}\alpha_{\text{EI},m}(\theta) = &\ (f_{\min,m} - \mu_m(\theta))\,\Phi\left(\zeta_m(\theta)\right) \\ &+ \sigma_m(\theta)\varphi\left(\zeta_m(\theta)\right),\end{aligned} \tag{10}$$

where $\zeta_m(\theta) = \sigma_m(\theta)^{-1}\left(f_{\min,m} - \mu_m(\theta)\right)$, and Φ and φ denote the cumulative distribution and probability density functions of the standard normal distribution, respectively. The subsequent query point θ_{m+1} is derived by maximising the acquisition function, i.e., we have

$$\theta_{m+1} = \text{argmax}_{\theta\in\Theta}\alpha_{\text{EI},m}(\theta), \tag{11}$$

which is solved using a global optimisation algorithm [31] to efficiently search the parameter space.

Note that, in learning the surrogate model f_m through GPR, it is assumed that $f \sim \mathcal{GP}(\mu, k)$, where $\mu : \Theta \to \mathbb{R}$ is the mean function, typically set to zero initially, and $k : \Theta \times \Theta \to \mathbb{R}$ is the covariance kernel. We employ here the *Squared Exponential* kernel defined as

$$k(\theta^{(1)}, \theta^{(2)}) = \sigma_f^2 \exp\left(-\frac{1}{2l^2}\|\theta^{(1)} - \theta^{(2)}\|^2\right), \tag{12}$$

with signal variance σ_f^2 controlling function variability and length scale l determining smoothness (larger l yields smoother, more global structure, while smaller l allows sharper local variations). The hyperparameters $\eta = [l, \sigma_f^2, \sigma_{\text{n}}^2]^\top$, including the noise variance σ_{n}^2, are estimated by maximising the log marginal likelihood [32] offline during GP training, defined as

$$\log p(y \mid \Theta, \eta) = -\frac{1}{2}y^\top K^{-1}y - \frac{1}{2}\text{logdet}(K) - \frac{n}{2}\log 2\pi, \tag{13}$$

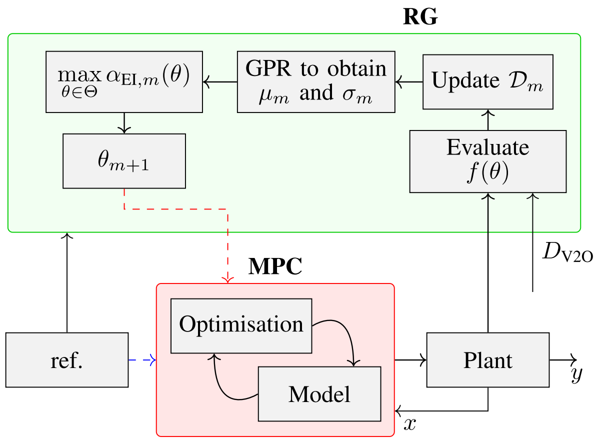

Fig. 3. Schematic of Bayesian Optimisation loop interacting with MPC for path following and obstacle avoidance. In the offline training phase only the RG reference is fed to the MPC depicted by a red dashed line. During online operation, the original reference in blue is fed until obstacle detection, prompting the RG augmentation.

Algorithm 1: BORG – GP surrogate pre-training via BO.

INPUT: Decision variable space Θ, objective function $f(\theta)$, max evaluations $m_{\max}$
Select initial configuration $\theta_0 \in \Theta$
Evaluate initial objective $y_0 = f(\theta_0)$ via system simulation
$\theta^* \leftarrow \theta_0$, $y^* \leftarrow y_0$, $S_0 \leftarrow \{\theta_0, y_0\}$
for $m = 1$ to $m_{\max}$ **do**
  Select new configuration $\theta_m \in \Theta$ by optimising acquisition function α_m
  $\theta_m \leftarrow \mathrm{argmax}_{\theta\in\Theta}\alpha_m(\theta; S_t)$
  Evaluate $f(\theta_m)$ to obtain y_m
  $S_m \leftarrow S_{m-1} \cup \{\theta_m, y_m\}$
  Update surrogate model and $\mathcal{D}_m$
  if $y_m < y^*$ **then**
    $\theta^* \leftarrow \theta_m$, $y^* \leftarrow y_m$
  end if
end for
OUTPUT: θ^* and y^*

where K is the covariance matrix over observed samples and y the corresponding function evaluations. The hyperparameters η can be estimated through negative log likelihood minimisation.

Following the above discussion, BO is used to efficiently explore the parameter space Θ, relying on a GPR to model the unknown objective function f. At iteration $m \in \mathbb{N}$, the GP provides a predictive posterior distribution characterised by a mean μ_m and variance σ_m^2, which are used to compute the Expected Improvement acquisition function as (10) and the next candidate parameter vector is selected by maximising this acquisition. The overall parameter tuning process consists of the following steps:

1) Sample an initial set of parameter vectors $\{\theta_m\}_{i=1}^{m_0}$ using Latin Hypercube Sampling or random sampling.
2) Evaluate $f(\theta_m)$ for each θ_m by running the closed-loop system.
3) Fit a (non-parametric) GP surrogate model to the collected data.
4) Iteratively select new candidates by maximising the acquisition function and updating the dataset.

The discussed iterative optimisation generates a sequence of parameter vectors $(\theta_m)_{m\in\mathbb{N}}$ that efficiently converges to optimal sigmoid parameters θ^* which enhance controller performance while avoiding exhaustive search over the parameter space. The process is summarised in Algorithm 1. Also, the interaction between the Bayesian optimiser and the MPC system is illustrated in Fig. 3. The dataset size remains moderate (on the order of tens to hundreds of evaluations, extending to at most thousands across scenarios), ensuring that GP training and covariance matrix inversion remain tractable in the offline phase.

D. Online Reference Governed MPC

We design an online implementation of the BORG system, leveraging a given pre-trained GP model [33], to perform real-time augmentation of the reference trajectory based on contextual information. The proposed scheme improves efficiency further and enables adaptive reference adjustment without incurring the computational overhead of full BO during runtime. Thus, to facilitate the mentioned features, an offline training phase is conducted using the BO framework described in the previous section as outlined in Algorithm 1, where the closed-loop system is simulated across diverse manoeuvre contexts and sigmoid reference parameters $\theta := [\vartheta_1, \vartheta_2, \vartheta_3]^\top$. The resulting dataset $\mathcal{D}_N = \{(\theta_m, f(\theta_m))\}_{i=1}^N$ captures system responses in terms of tracking error, control effort, and obstacle clearance, and subsequently, it is used to fit a GP model.

In the online phase, as demonstrated in Fig. 3, the pre-trained GP acts as a surrogate for the objective function $f(\theta)$, enabling rapid evaluation of reference trajectory quality based on the current driving context, defined by longitudinal velocity v_x and relative obstacle position $(x_{\mathrm{obs}}, y_{\mathrm{obs}})$. These form the GP input, which estimates expected performance for different sigmoid parametrisations.

The RG selects the optimal θ^* by maximising the acquisition function using the GP posterior conditioned on the current context, as outlined in Algorithm 2.

Expected improvement balances exploration and exploitation and can be tuned to favour exploration by lowering the improvement threshold. The introduced hybrid offline-online approach enables fast context-aware reference adaptation with minimal computational burden and improved tracking.

V. Efficient Data-Driven RG Training

To evaluate the performance and robustness of the BORG system, an efficient experimental design is required to balance data quality and practical feasibility. This section presents three strategies for generating training data for the offline Bayesian Optimisation loop Algorithm 1 and the GP model used online by the RG. Each strategy trades off data volume, training efficiency, and feasibility. A practical approach is the third hybrid scheme, combining rapid development and robustness refinement from

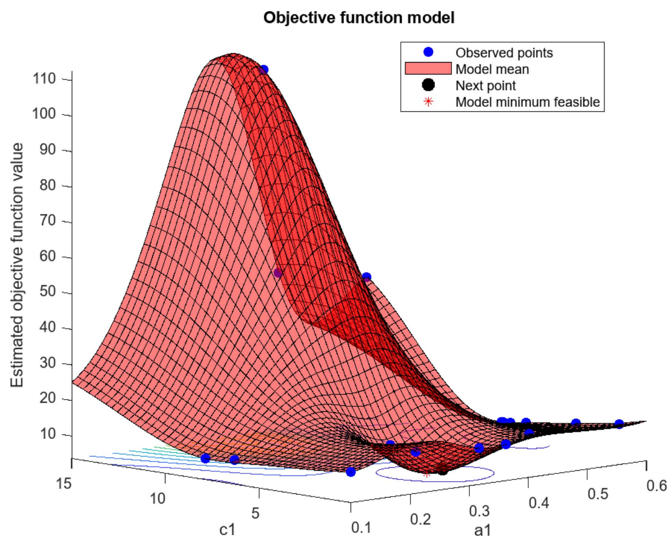

Fig. 4. BO objective landscape over sigmoid parameters a_1 and c_1.

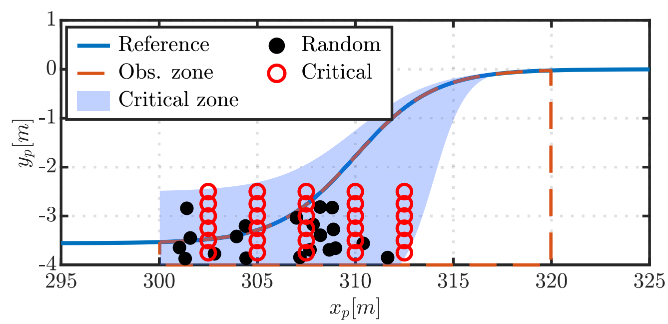

Fig. 5. Sigmoid with obstacle zones and random obstacle locations.

Algorithm 2: Online BORG – optimising unknown function f.

INPUT: Target function $f(\theta)$, training dataset D_N, number of iterations T
Pre-train a Gaussian Process:
$\mathcal{GP}(\hat{\mu}, \hat{k} \oplus \hat{\sigma}^2) \leftarrow$ PRE-TRAIN(D_N)
Initialize: $D_f \leftarrow \emptyset$
for $t = 1$ to T **do**
  Select next query point θ_t by optimising acquisition function:
   $\theta_t \leftarrow \text{argmax}_{\theta \in \Theta} \alpha(\theta; \mathcal{GP}(\hat{\mu}, \hat{k} \oplus \hat{\sigma}^2 \mid D_f))$
  Evaluate function output: $y_t \leftarrow f(\theta_t)$
  $D_f \leftarrow D_f \cup \{(\theta_t, y_t)\}$
end for
OUTPUT: D_f, θ_t

specific high risk conditions. Fig. 5 illustrates the sigmoid reference trajectory and obstacle zones used.

The objective is to sufficiently explore the contextual parameter space, vehicle velocity v_x and obstacle position $(x_{\text{obs}}, y_{\text{obs}})$, so that the GP model generalises across relevant driving scenarios. Each training sample corresponds to a simulated manoeuvre

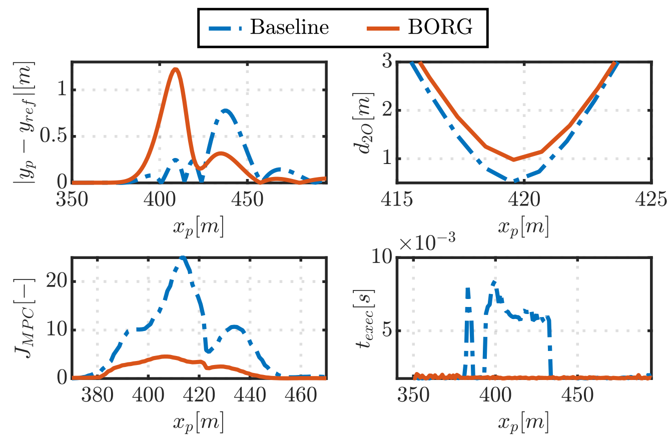

Fig. 6. Tracking error, minimum obstacle distance, control cost, and solver time.

with defined conditions, from which the performance cost $\Upsilon(\theta)$ is evaluated to update the GP.

1) Uniform random sampling: Sampling across a wide domain of vehicle velocities and obstacle locations promotes generalisation but requires extensive data, making it computationally expensive and impractical for physical collection, as many samples are low-risk and contribute little.

2) High-risk conditions: Sampling focuses on obstacles near the ego lane and higher speeds near handling limits. Resolution is finer in y_{obs}, which more strongly influences evasive actions, and coarser in x_{obs}. This enables a smaller yet informative dataset, though the resulting model may extrapolate poorly to less critical scenarios.

3) Critical-zone: This refines the second approach by sampling random obstacle positions within the critical zone while restricting velocity to two representative values: one for comfort and one near the handling limit. Thus, we lower the training effort while enabling interpolation between key speeds, assuming the optimal reference weakly depends on intermediate velocities; however, the simplified framework may lead to failing in rare edge cases. Training scenarios used velocities ranging between $55\,\text{km/h}$ and $80\,\text{km/h}$, with obstacles placed between $400\,\text{m}$ and $420\,\text{m}$ longitudinally and between $-4\,\text{m}$ and $-2\,\text{m}$ laterally.

VI. Numerical Experiments

We evaluate the proposed BORG against a nominal MPC with obstacle avoidance [12] in high-fidelity and uncertain environments, focusing on safety, control effort, and robustness.

A. Simulation Using High-Fidelity Vehicle Model

Validation is performed in IPG CarMaker using a high-fidelity SUV model with detailed suspension, steering, and powertrain dynamics. The baseline is a bicycle model MPC with Dugoff tyre forces [12], [26].

BORG is trained using 20 samples over critical obstacle positions ($380 \leq x_p \leq 425$ m, $-4.5 \leq y_p \leq -2$ m) at 50 and

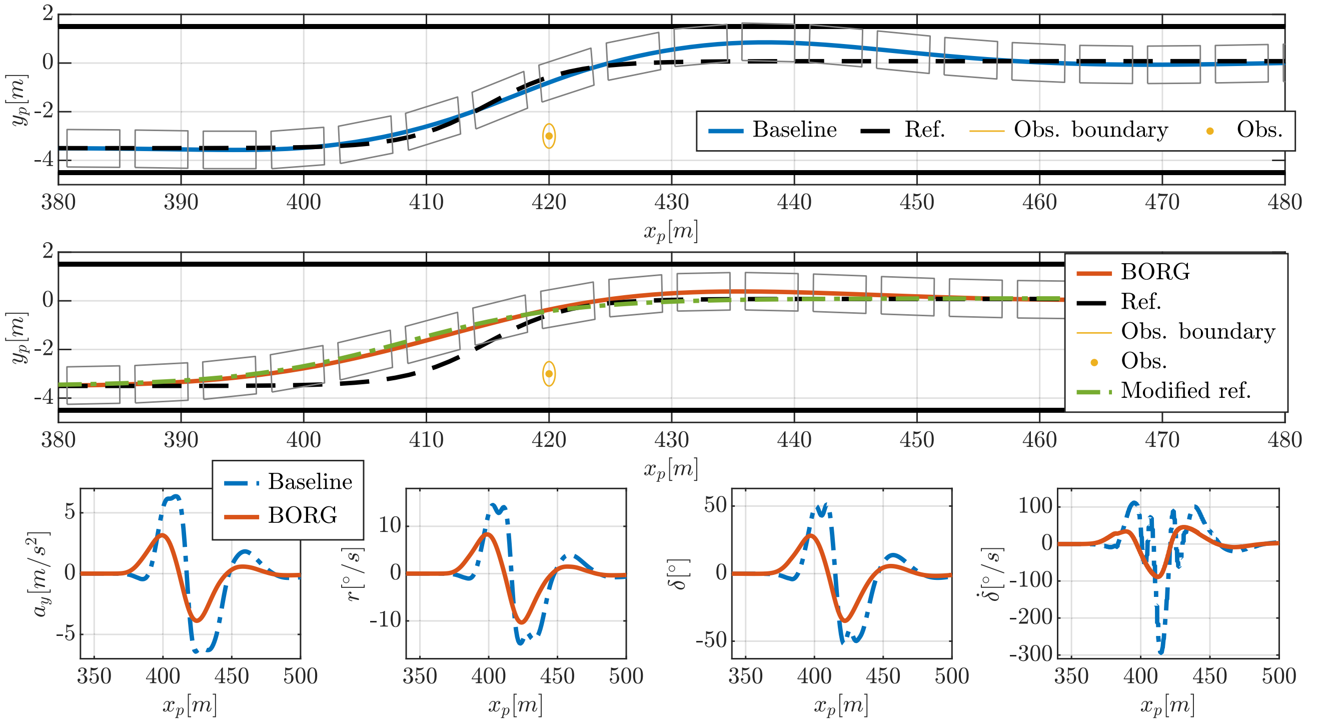

Fig. 7. Top: baseline MPC with obstacle avoidance. Middle: BORG with online reference adaptation. Bottom: lateral acceleration, yaw rate, steering angle, and steering rate.

TABLE I
PERFORMANCE AT 80 KM/H: RISE/SETTLING DISTANCE, OVERSHOOT, RMSE, AND MINIMUM OBSTACLE DISTANCE

	x_r	M_p	x_s	y_{RMSE}	$y^{\text{pre}}_{\text{RMSE}}$	$y^{\text{post}}_{\text{RMSE}}$	$\text{D2O}_{\min}$
Baseline	17.45	21.61%	102.62	0.23	0.10	0.28	0.52
BORG	26.47	8.61%	95.23	0.32	0.52	0.12	0.98

90 km/h, with BO used to train a GP for online reference adaptation.

As shown in Fig. 7, the baseline exhibits overshoot and road-edge violations. BORG increases safety margins by earlier deviation and improved post-obstacle convergence, yielding smoother responses (reduced lateral acceleration and yaw rate peaks), a $+89\,\%$ increase in minimum obstacle distance, and a $13\,\%$ overshoot reduction.

B. Robustness & Sensitivity Analysis

A trade-off is observed: pre-obstacle RMSE (before 420 m) increases due to earlier deviation, while post-obstacle performance improves, with reduced overshoot, shorter settling distance, and lower RMSE. Steering magnitude and jitter are also reduced, benefiting actuator usage and comfort (Table I).

BORG also reduces controller cost and solver time (Fig. 6), indicating that less aggressive manoeuvres avoid limit-handling conditions and keep the system in a more linear regime, supporting the use of simplified models with data-driven reference adaptation.

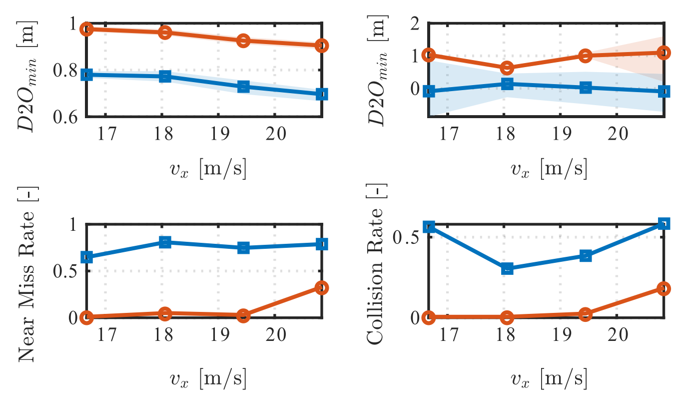

Fig. 8. Top: clearance under tyre mismatch. Middle: clearance under perception errors. Bottom: near-miss and collision rates (BORG vs. baseline).

The BO landscape (Fig. 4) shows high cost for aggressive manoeuvres with low clearance. The optimum reduces slope and shifts initiation earlier. Samples cluster near the optimum, indicating efficient convergence, although the flat minimum suggests potential for improved objective shaping.

Robustness is evaluated under perception noise and tyre model mismatch, which dominate over other variations. Ego localisation noise follows a Gaussian distribution centred at 0.1 m (tail 0.3 m), obstacle noise at 0.5 m (tail 0.1 m), and tyre dynamics vary by $\pm 20\%$ in Magic Formula parameters B_y and D_y [27].

As shown in Fig. 8, BORG consistently maintains larger obstacle clearance, with fewer near-miss and collision events, especially under perception noise. Although both methods degrade at higher speeds, BORG sustains safer distances across all velocities. This improvement stems from reduced manoeuvre aggressiveness, increasing vehicle-to-obstacle distance.

Overall, BORG enhances robustness by compensating for model mismatch and uncertainty through data-driven reference shaping. However, its reduced aggressiveness may underutilise vehicle capabilities in time-critical scenarios, and the one-shot reference update limits adaptability during execution.

VII. Conclusion

In this work, we have developed a novel control architecture, called BORG, that enhances obstacle avoidance and path following for automated vehicles by integrating a model-free data-driven reference governor within a non-linear model predictive framework. By optimising the parameters of a sigmoid-based trajectory modification offline using Bayesian optimisation, and adapting the reference in real-time using a Gaussian Process surrogate model, the proposed approach improves tracking performance and increases obstacle clearance. High-fidelity simulations demonstrated that BORG reduces overshoot, improves safety margins, and maintains lower control costs compared to the baseline while remaining computationally efficient. The model-free nature of the proposed method enables robustness to perception errors and model mismatch, addressing a key limitation of conventional planning-control pipelines. Although limited by its one-shot trajectory design and reduced aggressiveness, the proposed architecture effectively bridges the gap between motion planning and control, offering a promising direction for robust, real-time motion replanning.

References

[1] Y. K. Al-Nadawi, H. Al-Qassab, D. Kent, S. Pang, V. Srivastava, and H. Radha, "Design of robust path-following control system for self-driving vehicles using extended high-gain observer," in *Proc. Amer. Control Conf.*, 2020, pp. 1435–1440.

[2] F. Altché, P. Polack, and A. D. L. Fortelle, "High-speed trajectory planning for autonomous vehicles using a simple dynamic model," in *Proc. Int. Conf. Intell. Transp. Syst.*, 2017, pp. 1–7.

[3] A. Bertipaglia et al., "On the benefits of torque vectoring for automated collision avoidance at the limits of handling," *IEEE Trans. Veh. Technol.*, vol. 74, no. 6, pp. 8756–8771, Jun. 2025.

[4] D. Lenssen, A. Bertipaglia, F. Santafe, and B. Shyrokau, "Combined path following and vehicle stability control using model predictive control," SAE Int., Warrendale, PA, USA, Tech. Rep. 2023-01-0645, 2023.

[5] M. Werling, J. Ziegler, S. Kammel, and S. Thrun, "Optimal trajectory generation for dynamic street scenarios in a Frenét frame," in *Proc. IEEE Intell. Veh. Symp.*, 2010, pp. 987–993.

[6] D. González, J. Pérez, V. Milanés, and F. Nashashibi, "A review of motion planning techniques for automated vehicles," *IEEE Trans. Intell. Transp. Syst.*, vol. 17, no. 4, pp. 1135–1145, Apr. 2016.

[7] J. K. Subosits and J. C. Gerdes, "From the racetrack to the road: Real-time trajectory replanning for autonomous driving," *IEEE Trans. Intell. Veh.*, vol. 4, no. 2, pp. 309–320, Jun. 2019.

[8] A. Rupenyan, M. Khosravi, and J. Lygeros, "Performance-based trajectory optimization for path following control using Bayesian optimization," in *Proc. Conf. Decis. Control*, 2021, pp. 2116–2121.

[9] J. Ziegler, P. Bender, T. Dang, and C. Stiller, "Trajectory planning for Bertha–A local, continuous method," in *Proc. IEEE Intell. Veh. Symp.*, 2014, pp. 450–457.

[10] A. Liniger, A. Domahidi, and M. Morari, "Optimization-based autonomous racing of 1:43 scale RC cars," *Optimal Control Appl. Methods*, vol. 36, no. 5, pp. 628–647, 2015.

[11] B. Paden, M. Čáp, S. Z. Yong, D. Yershov, and E. Frazzoli, "A survey of motion planning and control techniques for self-driving urban vehicles," *IEEE Trans. Intell. Veh.*, vol. 1, no. 1, pp. 33–55, Mar. 2016.

[12] A. Bertipaglia, M. Alirezaei, R. Happee, and B. Shyrokau, "A learning-based model predictive contouring control for vehicle evasive manoeuvres," in *Proc. Int. Symp. Adv. Veh. Control*, 2024, pp. 632–638.

[13] J. Kabzan, L. Hewing, A. Liniger, and M. N. Zeilinger, "Learning-based model predictive control for autonomous racing," *IEEE Robot. Autom. Lett.*, vol. 4, no. 4, pp. 3363–3370, Oct. 2019.

[14] M. Khosravi, A. Eichler, N. Schmid, R. S. Smith, and P. Heer, "Controller tuning by Bayesian optimization an application to a heat pump," in *Proc. Eur. Control Conf.*, 2019, pp. 1467–1472.

[15] M. Khosravi, C. Konig, M. Maier, R. S. Smith, J. Lygeros, and A. Rupenyan, "Safety-aware cascade controller tuning using constrained Bayesian optimization," *IEEE Trans. Ind. Electron.*, vol. 70, no. 2, pp. 2128–2138, Feb. 2023.

[16] G. Savaia, Y. Sohn, S. Formentin, G. Panzani, M. Corno, and S. M. Savaresi, "Experimental automatic calibration of a semi-active suspension controller via Bayesian optimization," *Control Eng. Pract.*, vol. 112, 2021, Art. no. 104826.

[17] E. Catenaro, L. Aarnoudse, S. Formentin, and T. Oomen, "Efficient tuning for motion control in diverse systems: A Bayesian framework," *IFAC-PapersOnLine*, vol. 58, no. 15, pp. 354–359, 2024.

[18] I. Kolmanovsky and E. Gilbert, "Theory and computation of disturbance invariant sets for discrete-time linear systems," *Math. Problems Eng.*, vol. 4, no. 4, pp. 317–367, 1998.

[19] I. Kolmanovsky, A. Bemporad, and E. Garone, "Vehicle systems with reference and command governors: Opportunities for semi-autonomous driving," in *Proc. Conf. Decis. Control*, Osaka, Japan, 2015, pp. 4012–4017.

[20] X. Li et al., "Lane keeping using Lyapunov function-based reference governor: An optimization-free approach," in *Proc. Amer. Control Conf.*, 2025, pp. 2342–2349.

[21] S. D. Cairano, T. Skibik, A. P. Vinod, A. Weiss, K. Berntorp, and Y. Okura, "Parametrized maneuvers governor for decision-making in automated driving," in *Proc. Nonlinear Constrained Control*, 2025, pp. 329–354.

[22] E. Garone, S. D. Cairano, and I. Kolmanovsky, "Reference and command governors for systems with constraints: A survey on theory and applications," *Automatica*, vol. 75, pp. 306–328, Jan. 2017.

[23] A. Carvalho, S. Lefévre, H. Gao, and F. Borrelli, "Stochastic predictive control of autonomous vehicles in uncertain environments," in *Proc. Int. Symp. Adv. Veh. Control*, Tokyo, Japan, 2014, pp. 712–717.

[24] M. Khosravi, V. Behrunani, R. S. Smith, A. Rupenyan, and J. Lygeros, "Cascade control: Data-driven tuning approach based on Bayesian optimization," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 382–387, 2020.

[25] M. Khosravi, V. N. Behrunani, P. Myszkorowski, R. S. Smith, A. Rupenyan, and J. Lygeros, "Performance-driven cascade controller tuning with Bayesian optimization," *IEEE Trans. Ind. Electron.*, vol. 69, no. 1, pp. 1032–1042, 2021.

[26] N. Chowdhri, L. Ferranti, F. S. Iribarren, and B. Shyrokau, "Integrated nonlinear model predictive control for automated driving," *Control Eng. Pract.*, vol. 106, 2021, Art. no. 104654.

[27] H. B. Pacejka and E. Bakker, "The magic formula tyre model," *Veh. System Dyn.*, vol. 21, no. S1, pp. 1–18, 1992.

[28] N. Ding and S. Taheri, "A modified Dugoff tire model for combined-slip forces," *Tire Sci. Technol.*, vol. 38, no. 3, pp. 228–244, 2010.

[29] I. Bae, J. H. Kim, and S. Kim, "Steering rate controller based on curvature of trajectory for autonomous driving vehicles," in *Proc. IEEE Intell. Veh. Symp.*, 2013, pp. 1381–1386.

[30] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA, USA: MIT Press, 2006.

[31] A. Lambora, K. Gupta, and K. Chopra, "Genetic algorithm–A literature review," in *Proc. Int. Conf. Mach. Learn., Big Data, Cloud Parallel Comput.*, 2019, pp. 380–384.

[32] J. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl, "Algorithms for hyper-parameter optimization," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 24, 2011, pp. 2546–2554.

[33] Z. Wang et al., "Pre-trained Gaussian processes for Bayesian optimization," *J. Mach. Learn. Res.*, vol. 25, no. 212, pp. 1–83, 2024.